

Curriculum Vitae

Ethan L. Miller

October 2008

Computer Science Department
University of California, Santa Cruz
1156 High Street, MS SOE3
Santa Cruz, CA 95064

PHONE: +1 (831) 345-4864
FAX: +1 (815) 550-1178
EMAIL: elm@cs.ucsc.edu
<http://www.cs.ucsc.edu/~elm/>

EMPLOYMENT HISTORY

2000– Computer Science Department, University of California, Santa Cruz
 Professor (2008–)
 Associate Professor (2002–2008)
 Assistant Professor (2000–2002)

1999 System Architect, Endeca, Cambridge, MA

1994–2000 Assistant Professor, Computer Science and Electrical Engineering Department, University of Maryland Baltimore County

1988–1994 Research Assistant, Computer Science Division, University of California at Berkeley

1988–1990 Teaching Associate, Computer Science Division, University of California at Berkeley

1987–1988 Software Engineer, BBN Laboratories, Cambridge, MA

1986 Summer intern, GTE Government Systems, Rockville, MD

EDUCATION

1995 Ph. D., University of California at Berkeley, Computer Science (advisor: Randy Katz)
 Thesis: *Storage Hierarchy Management for Scientific Computing*

1990 M. S., University of California at Berkeley, Computer Science

1987 Sc. B., Brown University, Computer Science, *magna cum laude*

CONSULTING

2008 Consulting expert witness, Shepperd Mullin

2007 Hewlett Packard Laboratories

2006–2007 Consulting and testifying expert witness, Fish & Richardson

2005–2006 Veritas (now Symantec) Corporation

2003 Hewlett Packard Laboratories

2001–2005 Expert witness, Bartlit, Beck, Herman, Palenchar & Scott

2000 Expert witness, Fish & Richardson

1998–2001 Expert witness, Hopgood, Calimafde, Judlowe & Mondolino

1998 Web site architect, Ambleside Logic

HONORS

- 2005 Best Long Paper award, StorageSS, 2005.
- 2004 Best Paper award, MASCOTS 2004.
- 2001 Elevated to Senior Member, IEEE.
- 1987 William Gaston Prize for Academic Excellence (award made to top graduating students at Brown University).
- 1987 Elected to Sigma Xi, Brown University.

RESEARCH FUNDING

In addition to the grants listed below, I have helped bring in industrial funding for the Storage Systems Research Center, which, since 2002, has received \$100,000–\$200,000 per year from companies including Data Domain, Hewlett Packard, Hitachi, IBM, LSI, NetApp, Seagate, and Symantec.

Grants

- 2006–2011 Co-PI, *Petascale Data Storage Institute*, \$11,250,000, Department of Energy (lead PI: Garth Gibson; UCSC PI: Darrell Long; other UC Santa Cruz co-PI is Scott Brandt). The Institute funds research at three universities and five national laboratories, with \$1,500,000 going to UC Santa Cruz.
- 2006–2009 Co-PI, *File System Tracing, Replaying, Profiling, and Analysis on HEC Systems*, National Science Foundation, \$760,252 (PI: Erez Zadok, Stony Brook University; co-PI: Klaus Mueller, Stony Brook University).
- 2006–2007 Co-PI, *Institute for Scalable Scientific Data Management*, Los Alamos National Laboratory, \$750,000 (PI: Darrell Long; co-PI: Scott Brandt; co-PI: Carlos Maltzahn).
- 2005–2006 PI, *Adaptive Workload-Aware Algorithms for Heterogeneous Storage Systems*, UC MICRO, \$80,269 (includes \$45,000 gift from Veritas).
- 2005–2006 Co-PI, *Scalable File Systems For High Performance Computing*, Lawrence Livermore National Laboratory, Los Alamos National Laboratory, and Sandia National Laboratory, \$250,000 (PI: Scott Brandt; co-PI: Darrell Long; co-PI: Martín Abadi; co-PI: Carlos Maltzahn).
- 2005–2006 Co-PI, *Institute for Scalable Scientific Data Management*, Los Alamos National Laboratory, \$750,000 (PI: Darrell Long; co-PI: Scott Brandt).
- 2005 Co-PI, *Trustworthy Computing Curriculum Development*, Microsoft, \$50,000 (PI: Ira Pohl; co-PI: Martín Abadi; co-PI: Jim Whitehead).
- 2004 PI, *Research in Storage and Networks*, Hewlett-Packard Laboratories, \$25,000.
- 2003 PI, *Research in Storage and Networks*, Hewlett-Packard Laboratories, \$35,000.
- 2003–2006 PI, *Building High-performance, Reliable Storage Systems Using Magnetic RAM*, National Science Foundation, \$414,000 (co-PI: Scott Brandt).
- 2002–2005 Co-PI, *Scalable File Systems for High Performance Computing*, Department of Energy, \$900,000 (PI: Darrell Long; co-PI: Scott Brandt; co-PI: Katia Obraczka).
- 2002 PI, *Research in Storage and Networks*, Hewlett-Packard Laboratories, \$42,000.
- 2001–2002 Co-PI, *Building a High-Performance Storage System from Commodity Components*, Lawrence Livermore National Laboratory, \$65,000 (PI: Darrell Long; co-PI: Scott Brandt).
- 2001 PI, *Research in Storage and Networks*, Hewlett-Packard Laboratories, \$38,000.
- 1998–2001 Faculty Fellowship, University of Maryland Institute for Advanced Computer Studies, \$30,000.

- 1997–2001 Co-PI, *Center for Architectures for Data-Driven Information Processing*, Department of Defense, \$3,000,000 (PI: Charles Nicholas; co-PI: David Ebert).
- 1997–1998 PI, *Scalable Benchmarks for Mass Storage Systems*, NASA Ames Research Center, \$200,000.
- 1995–1997 Co-PI, *Scalability of the TELLTALE Dynamic Hypertext Environment*, Department of Defense, \$100,000 (PI: Charles Nicholas).
- 1995–1998 Faculty Fellowship, University of Maryland Institute for Advanced Computer Studies, \$30,000.
- 1989–1993 Graduate Fellowship, National Science Foundation.
- 1988–1989 Cal MICRO Fellowship, University of California at Berkeley.

PUBLICATIONS

Journals

- J11. Andrew W. Leung, Minglong Shao, Tim Bisson, Shankar Pasupathy, and **Ethan L. Miller**, “High-Performance Metadata Indexing and Search in Petascale Data Storage Systems,” *Journal of Physics: Conference Series* **125** (2008) 012069, July 2008.
- J10. Carlos Maltzahn, Nikhil Bobb, Mark W. Storer, Damian Eads, Scott A. Brandt, and **Ethan L. Miller**, “Graffiti: A Framework for Testing Collaborative Distributed Metadata,” *Distributed Data & Structures* **7**, *Proceedings in Informatics* **21**, Carleton Scientific, 2007, pages 97–111. Extended version of the WDAS 2006 workshop paper.
- J9. Bo Hong, Scott A. Brandt, Darrell D. E. Long, **Ethan L. Miller**, and Ying Lin, “Using MEMS-Based Storage in Computer Systems—MEMS Storage Device Modeling and Management,” *ACM Transactions on Storage* **2**(2), May 2006, pages 139–160.
- J8. Qin Xin, Thomas J. E. Schwarz, and **Ethan L. Miller**, “Availability in Global Peer-To-Peer Storage Systems,” *Distributed Data & Structures* **6**, *Proceedings in Informatics* **20**, Carleton Scientific, 2005, pages 63–77. Extended version of the WDAS 2004 workshop paper.
- J7. Ismail Ari, Ahmed Amer, Robert Gramacy, **Ethan L. Miller**, Scott A. Brandt, and Darrell D. E. Long, “ACME: Adaptive Caching Using Multiple Experts,” *Distributed Data & Structures* **5**, *Proceedings in Informatics* **14**, Carleton Scientific, 2002, pages 143–158. Extended version of the WDAS 2002 workshop paper.
- J6. **Ethan Miller**, Dan Shen, Junli Liu, and Charles Nicholas, “Performance and Scalability of a Large-Scale N-gram Based Information Retrieval System,” *Journal of Digital Information* **1**(5), January 2000, 25 pages (online refereed journal).
- J5. Christopher Shaw, James Kukla, Ian Soboroff, David Ebert, Charles Nicholas, Amen Zwa, **Ethan Miller**, and D. Aaron Roberts, “Interactive Volumetric Information Visualization for Document Corpus Management,” *International Journal on Digital Libraries* **2**(2–3), 1999, pages 144–156.
- J4. Jeffrey Hollingsworth, **Ethan Miller**, and Kennedy Akala, “Binary Version Management for Computational Grids,” *Parallel Processing Letters* **9**(2), June 1999, pages 215–225.
- J3. **Ethan Miller** and Randy Katz, “RAMA: An Easy-To-Use, High-Performance Parallel File System,” *Parallel Computing* **23**(4), July 1997, pages 419–446.
- J2. David Ebert, Amen Zwa, **Ethan Miller**, Chris D. Shaw, and D. Aaron Roberts, “Two-handed Volumetric Document Corpus Management,” *IEEE Computer Graphics and Applications* **17**(4), July 1997, pages 60–62.

- J1. Peter Chen, Edward Lee, Ann Drapeau, Ken Lutz, **Ethan Miller**, Srinu Seshan, Ken Shirriff, David Patterson, and Randy Katz, "Performance and Design Evaluation of the RAID-II Storage Server," *Journal of Distributed and Parallel Databases* 2(3), July 1994, pages 243–260.

Refereed Conference & Workshop Papers

- C84. Kevin M. Greenan, Darrell D. E. Long, **Ethan L. Miller**, Thomas J. Schwarz, S. J., and Jay J. Wylie, "A Spin-Up Saved is Energy Earned: Achieving Power-Efficient, Erasure-Coded Storage" *Proceedings of the 4th Workshop on Hot Topics in System Dependability (HotDep'08)*, San Diego, CA: USENIX, December 2008.
- C83. Mark W. Storer, Kevin M. Greenan, Ian Adams, **Ethan L. Miller**, Darrell D. E. Long, and Kaladhar Voruganti, "Logan: Automatic Management for Evolvable, Large-Scale, Archival Storage," *Proceedings of the 3rd International Workshop on Petascale Data Storage (PDSW08)*, held in conjunction with SC2008, Austin, TX, November 2008.
- C82. Andrew W. Leung and **Ethan L. Miller**, "Scalable Full-Text Search for Petascale File Systems," *Proceedings of the 3rd International Workshop on Petascale Data Storage (PDSW08)*, held in conjunction with SC2008, Austin, TX, November 2008.
- C81. Mark W. Storer, Kevin Greenan, Darrell D. E. Long, and **Ethan L. Miller**, "Secure Data Deduplication," *Proceedings of the 4th Workshop on Storage Security and Survivability (StorageSS 2008)*, held in conjunction with the 15th ACM Conference on Computer and Communications Security (CCS 2008), Alexandria, VA, October 2008.
- C80. Mohammed G. Khatib, **Ethan L. Miller** and Pieter H. Hartel, "Workload-Based Configuration of MEMS-Based Storage Devices for Mobile Systems," *Proceedings of the 8th ACM & IEEE Conference on Embedded Software (EMSOFT '08)*, Atlanta, GA: ACM / IEEE, October 2008.
- C79. Kevin M. Greenan, **Ethan L. Miller**, and Thomas Schwarz, "Optimizing Galois Field Arithmetic for Diverse Processor Architectures," *Proceedings of the 16th IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2008)*, Baltimore, MD: IEEE, September 2008, 10 pages (CD-ROM).
- C78. Andrew W. Leung, Shankar Pasupathy, Garth Goodson, and **Ethan L. Miller**, "Measurement and Analysis of Large-Scale Enterprise Network File System Workloads," *Proceedings of the 2008 USENIX Annual Technical Conference (USENIX '08)*, Boston, MA: USENIX, June 2008, pages 213–226.
- C77. Kevin M. Greenan, **Ethan L. Miller**, and Jay J. Wylie, "Reliability of XOR-based erasure codes on heterogeneous devices," *Proceedings of the 38th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN 2008)*, Anchorage, AK: IEEE, June 2008, pages 147–156.
- C76. Neerja Bhatnagar, Kevin M. Greenan, Rosie Wacha, **Ethan L. Miller**, and Darrell D. E. Long, "Energy-Reliability Tradeoffs in Sensor Networks," *Proceedings of the 5th Workshop on Embedded Networked Sensors (HotEmNets 2008)*, Charlottesville, VA: ACM, June 2008.
- C75. Mark W. Storer, Kevin M. Greenan, **Ethan L. Miller**, and Kaladhar Voruganti, "Pergamum: Replacing Tape with Energy Efficient, Reliable, Disk-Based Archival Storage," *Proceedings of the 6th Conference on File and Storage Technologies (FAST '08)*, San Jose, CA: USENIX, February 2008, pages 1–16.

- C74. Jonathan Koren, Yi Zhang, Sasha Ames, Andrew W. Leung, Carlos Maltzahn, and **Ethan Miller**, “Searching and Navigating Petabyte Scale File Systems Based on Facets,” *Proceedings of the 2nd International Workshop on Petascale Data Storage (PDSW07)*, held in conjunction with *SC2007*, Reno, NV, November 2007, pages 21–25.
- C73. Andrew Leung, **Ethan L. Miller**, and Stephanie Jones, “Scalable Security for Petascale Parallel File Systems,” *Proceedings of SC2007*, Reno, NV: ACM, November 2007.
- C72. Kevin M. Greenan, **Ethan L. Miller**, Thomas J. E. Schwarz, S.J., and Darrell D. E. Long, “Disaster Recovery Codes: Increasing Reliability with Large-Stripe Error Correction Codes,” *Proceedings of the 3rd International Workshop on Storage Security and Survivability (StorageSS 2007)*, held in conjunction with the *14th ACM Conference on Computer and Communications Security (CCS 2007)*, Alexandria, VA, October 2007.
- C71. Neerja Bhatnagar and **Ethan L. Miller**, “A Secure and Reliable File System for Sensor Nodes,” *Proceedings of the 3rd International Workshop on Storage Security and Survivability (StorageSS 2007)*, held in conjunction with the *14th ACM Conference on Computer and Communications Security (CCS 2007)*, Alexandria, VA, October 2007.
- C70. Kevin Greenan and **Ethan L. Miller**, “PRIMS: Making NVRAM Suitable for Extremely Reliable Storage,” *Proceedings of the 3rd Workshop on Hot Topics in System Dependability (HotDep’07)*, Edinburgh, UK: USENIX, June 2007, 4 pages.
- C69. Mark W. Storer, Kevin Greenan, **Ethan L. Miller**, and Kaladhar Voruganti, “POTSHARDS: Secure Long-Term Storage Without Encryption,” *Proceedings of the 2007 USENIX Annual Technical Conference (USENIX ’07)*, Santa Clara, CA: USENIX, June 2007, pages 143–156.
- C68. Sage A. Weil, Scott A. Brandt, **Ethan L. Miller**, and Carlos Maltzahn, “CRUSH: Controlled, Scalable And Decentralized Placement Of Replicated Data,” *Proceedings of SC2006*, Tampa, FL: ACM, November 2006, 12 pages (published on CD-ROM).
- C67. Sage A. Weil, Scott A. Brandt, **Ethan L. Miller**, Darrell D. E. Long, and Carlos Maltzahn, “Ceph: A Scalable, High-Performance Distributed File System,” *Proceedings of the 7th Conference on Operating Systems Design and Implementation (OSDI ’06)*, Seattle, WA, November 2006, pages 307–320.
- C66. Mark W. Storer, Kevin M. Greenan, and **Ethan L. Miller**, “Long-Term Threats to Secure Archives,” *Proceedings of the 2nd International Workshop on Storage Security and Survivability (StorageSS 2006)*, held in conjunction with the *13th ACM Conference on Computer and Communications Security (CCS 2006)*, Virginia, October 2006, pages 9–16.
- C65. Andrew Leung and **Ethan L. Miller**, “Scalable Security for Large, High Performance Storage Systems,” *Proceedings of the 2nd International Workshop on Storage Security and Survivability (StorageSS 2006)*, held in conjunction with the *13th ACM Conference on Computer and Communications Security (CCS 2006)*, Virginia, October 2006, pages 29–40.
- C64. Kevin M. Greenan and **Ethan L. Miller**, “Reliability Mechanisms for File Systems Using Non-Volatile Memory as a Metadata Store,” *Proceedings of the 6th ACM & IEEE Conference on Embedded Software (EMSOFT ’06)*, Seoul, Korea, October 2006, pages 178–187.
- C63. Deepavali Bhagwat, Kristal Pollack, Darrell D. E. Long, **Ethan L. Miller**, Jehan-François Pâris, and Thomas Schwarz, S. J., “Providing High Reliability in a Minimum Redundancy Archival Storage System,” *Proceedings of the 14th IEEE/ACM International Symposium on Modeling, Analysis, and*

- Simulation of Computer and Telecommunication Systems (MASCOTS 2006)*, Monterey, CA, September 2006, pages 413–421.
- C62. Thomas J. E. Schwarz and **Ethan L. Miller**, “Store, Forget, and Check: Using Algebraic Signatures to Check Remotely Administered Storage,” *Proceedings of the 26th International Conference on Distributed Computing Systems (ICDCS ’06)*, Lisboa, Portugal, July 2006, 10 pages (published on CD-ROM).
- C61. Sasha Ames, Nikhil Bobb, Kevin M. Greenan, Owen S. Hofmann, Mark W. Storer, Carlos Maltzahn, **Ethan L. Miller**, and Scott A. Brandt, “LiFS: An Attribute-Rich File System for Storage Class Memories,” *Proceedings of the 23rd IEEE / 14th NASA Goddard Conference on Mass Storage Systems and Technologies (MSST 2006)*, College Park, MD, May 2006, pages 63–76.
- C60. Nikhil Bobb, Damian Eads, Mark W. Storer, Scott A. Brandt, Carlos Maltzahn, and **Ethan L. Miller**, “Graffiti: A Framework for Testing Collaborative Distributed Metadata,” *7th Workshop on Distributed Data and Structures (WDAS 2006)*, Santa Clara, CA, January 2006.
- C59. Mark Storer, Kevin Greenan, **Ethan L. Miller**, and Carlos Maltzahn, “POTSHARDS: Storing Data for the Long-term Without Encryption,” *Proceedings of the 3rd International IEEE Security in Storage Workshop*, San Francisco, CA, December 2005.
- C58. Christopher Olson and **Ethan L. Miller**, “Secure Capabilities for a Petabyte-Scale Object-Based Distributed File System,” *Proceedings of the International Workshop on Storage Security and Survivability (StorageSS)*, held in conjunction with the *12th ACM Conference on Computer and Communications Security (CCS 2005)*, Fairfax, VA, November 2005, pages 64–73. Received Best Paper award.
- C57. Qin Xin, Thomas J. E. Schwarz, S. J., and **Ethan L. Miller**, “Disk Infant Mortality in Large Storage Systems,” *Proceedings of the 13th IEEE/ACM International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2005)*, Atlanta, GA, September 2005, pages 125–134.
- C56. Qin Xin, **Ethan L. Miller**, Thomas J. E. Schwarz, S. J., and Darrell D. E. Long, “Impact Of Failure On Interconnection Networks for Large Storage Systems,” *Proceedings of the 22nd IEEE / 13th NASA Goddard Conference on Mass Storage Systems and Technologies (MSST 2005)*, Monterey, CA, April 2005, pages 189–196.
- C55. Alexander Ames, Nikhil Bobb, Scott A. Brandt, Adam Hiatt, Carlos Maltzahn, **Ethan L. Miller**, Alisa Neeman, and Deepa Tuteja, “Richer File System Metadata Using Links and Attributes,” *Proceedings of the 22nd IEEE / 13th NASA Goddard Conference on Mass Storage Systems and Technologies (MSST 2005)*, Monterey, CA, April 2005, pages 49–60.
- C54. Sage A. Weil, Kristal T. Pollack, Scott A. Brandt, and **Ethan L. Miller**, “Dynamic Metadata Management for Petabyte-scale File Systems,” *Proceedings of SC2004*, Pittsburgh, PA: ACM, November 2004. Nominee, Best Student Paper.
- C53. Thomas J. E. Schwarz, S. J., Qin Xin, **Ethan L. Miller**, Darrell D. E. Long, Andy Hospodor, and Spencer Ng, “Disk Scrubbing in Large Archival Storage Systems,” *Proceedings of the 12th IEEE/ACM International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2004)*, Volendam, Netherlands, October 2004, pages 409–418. Received Best Paper award.

- C52. Nathan K. Edel, Deepa Tuteja, **Ethan L. Miller**, and Scott A. Brandt, “MRAMFS: A Compressing File System for Non-Volatile RAM,” *Proceedings of the 12th IEEE/ACM International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2004)*, Volendam, Netherlands, October 2004, pages 596–603.
- C51. Thomas J. E. Schwarz, Qin Xin, and **Ethan L. Miller**, “Availability in Global Peer-To-Peer Storage Systems,” *6th Workshop on Distributed Data and Structures (WDAS 2004)*, Lausanne, Switzerland, July 2004.
- C50. Nathan K. Edel, **Ethan L. Miller**, Karl S. Brandt, and Scott A. Brandt, “Measuring the Compressibility of Metadata and Small Files for Disk/NVRAM Hybrid Storage Systems,” *Proceedings of the 2004 International Symposium on Performance Evaluation of Computer and Telecommunication Systems (SPECTS’04)*, San Jose, CA, July 2004.
- C49. Ismail Ari and **Ethan L. Miller**, “Caching Support for Push-Pull Data Dissemination using Data-snooping Routers,” *Proceedings of the 10th International Conference on Parallel and Distributed Systems (ICPADS)*, Newport Beach, CA: IEEE, July 2004, pages 101–108.
- C48. Qin Xin, **Ethan L. Miller**, and Thomas J. E. Schwarz, S. J. “Evaluation of Distributed Recovery in Large-Scale Storage Systems,” *Proceedings of the 13th IEEE International Symposium on High Performance Distributed Computing (HPDC-13)*, Honolulu, HI: IEEE, June 2004, pages 172–181.
- C47. Sage A. Weil, Scott A. Brandt, **Ethan L. Miller**, and Kristal T. Pollack, “Intelligent Metadata Management for a Petabyte-Scale File System,” 2nd Intelligent Storage Workshop, University of Minnesota, May 2004.
- C46. Emilia Rosti and **Ethan L. Miller**, “Security Threats and Responses for Object-Based Storage Devices,” 2nd Intelligent Storage Workshop, University of Minnesota, May 2004 (poster presentation).
- C45. R. J. Honicky and **Ethan L. Miller**, “Replication Under Scalable Hashing: A Family of Algorithms for Scalable Decentralized Data Distribution,” *Proceedings of the 18th International Parallel and Distributed Processing Symposium (IPDPS 2004)*, Santa Fe, NM: IEEE, April 2004, 10 pages (published on CD-ROM).
- C44. Feng Wang, Scott A. Brandt, **Ethan L. Miller**, and Darrell D. E. Long, “OBFS: A File System for Object-Based Storage Devices,” *Proceedings of the 21st IEEE / 12th NASA Goddard Conference on Mass Storage Systems and Technologies (MSST 2004)*, College Park, MD, April 2004, pages 283–300.
- C43. Andy Hospodor and **Ethan L. Miller**, “Interconnection Architectures for Petabyte-Scale High-Performance Storage Systems,” *Proceedings of the 21st IEEE / 12th NASA Goddard Conference on Mass Storage Systems and Technologies (MSST 2004)*, College Park, MD, April 2004, pages 273–281.
- C42. Feng Wang, Qin Xin, Bo Hong, Scott A. Brandt, **Ethan L. Miller**, Darrell D. E. Long, and Tyce T. McLarty, “File System Workload Analysis for Large Scale Scientific Computing Applications,” *Proceedings of the 21st IEEE / 12th NASA Goddard Conference on Mass Storage Systems and Technologies (MSST 2004)*, College Park, MD, April 2004, pages 139–152.
- C41. Bo Hong, Scott A. Brandt, Darrell D. E. Long, **Ethan L. Miller**, Karen A. Glocer, and Zachary N. J. Peterson, “Zone-Based Shortest Positioning Time First Scheduling for MEMS-Based Storage Devices,” *Proceedings of the 11th IEEE/ACM International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2003)*, Orlando, FL, October 2003, pages 104–113.

- C40. Ismail Ari, Bo Hong, **Ethan L. Miller**, Scott A. Brandt and Darrell D. E. Long, “Managing Flash Crowds On The Internet,” *Proceedings of the 11th IEEE/ACM International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2003)*, Orlando, FL, October 2003, pages 246–249.
- C39. R. J. Honicky and **Ethan L. Miller**, “A Fast Algorithm for Online Placement and Reorganization of Replicated Data,” *Proceedings of the 17th International Parallel and Distributed Processing Symposium*, Nice, France: IEEE, April 2003, 10 pages (published on CD-ROM). Also available as Technical Report UCSC-CRL-02-36.
- C38. Qin Xin, **Ethan L. Miller**, Thomas Schwarz, Scott A. Brandt, Darrell D. E. Long, and Witold Litwin, “Reliability Mechanisms for Very Large Storage Systems,” *Proceedings of the 20th IEEE / 11th NASA Goddard Conference on Mass Storage Systems and Technologies (MSST 2003)*, San Diego, CA: IEEE, April 2003, pages 146–156.
- C37. Scott A. Brandt, **Ethan L. Miller**, Darrell D. E. Long, and Lan Xue, “Efficient Metadata Management in Large Distributed File Systems,” *Proceedings of the 20th IEEE / 11th NASA Goddard Conference on Mass Storage Systems and Technologies (MSST 2003)*, San Diego, CA: IEEE, April 2003, pages 290–298.
- C36. Ying Lin, Scott A. Brandt, Darrell D. E. Long, and **Ethan L. Miller**, “Power Conservation Strategies for MEMS-based Storage Devices,” *Proceedings of the 10th International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2002)*, Fort Worth, TX: IEEE, October 2002, pages 53–62.
- C35. Scott A. Banachowski, Zachary N. J. Peterson, **Ethan L. Miller**, and Scott A. Brandt, “Intra-file Security for a Distributed File System,” *Proceedings of the 19th IEEE Symposium on Mass Storage Systems and Technologies*, College Park, MD: IEEE, April 2002, pages 153–163.
- C34. Ismail Ari, Ahmed Amer, **Ethan Miller**, Scott Brandt, and Darrell Long. “Who is more adaptive? ACME: adaptive caching using multiple experts,” *Workshop on Distributed Data and Structures (WDAS 2002)*, Paris, France, March 2002.
- C33. **Ethan L. Miller**, Darrell D. E. Long, William E. Freeman, and Benjamin C. Reed, “Strong Security for Network-Attached Storage,” *Proceedings of the First Conference on File and Storage Technologies (FAST)*, Monterey, CA: Usenix, January 2002, pages 1–13.
- C32. **Ethan L. Miller**, Scott A. Brandt and Darrell D. E. Long, “HeRMES: High-Performance Reliable MRAM-Enabled Storage,” *Proceedings of the 8th IEEE Workshop on Hot Topics in Operating Systems (HotOS-VIII)*, Elmau, Germany: IEEE, May 2001, pages 83–87.
- C31. Lee Butler, Travis Atkison, and **Ethan Miller**, “Comparing CPU Performance Between and Within Processor Families,” *Proceedings of the 25th Annual International Conference on Computer Measurement and Performance (CMG 2000)*, Orlando, FL, December 2000, pages 421–430.
- C30. **Ethan Miller** and Jon Squire, “esim: A Structural Design Language and Simulator for Computer Architecture Education,” *2000 Workshop on Computer Architecture Education (WCAE 2000)*, Vancouver, Canada: ACM & IEEE, June 2000, pages 42–48.
- C29. William Freeman and **Ethan Miller**, “Design for A Decentralized Security System For Network Attached Storage,” *Proceedings of the 8th Goddard Conference on Mass Storage Systems and Technologies / 17th IEEE Symposium on Mass Storage Systems*, College Park, MD, March 2000, pages 361–373.

- C28. Timothy Gibson and **Ethan Miller**, “An Improved Long-Term File Usage Prediction Algorithm,” *Proceedings of the 25th Annual International Conference on Computer Measurement and Performance (CMG '99)*, Reno, NV: CMG, December 1999, pages 639–648.
- C27. William E. Freeman and **Ethan L. Miller**, “An Experimental Analysis of Cryptographic Overhead in Performance-Critical Systems,” *Proceedings of the 7th International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS '99)*, College Park, MD: IEEE, October 1999, pages 348–357.
- C26. **Ethan Miller**, Dan Shen, Junli Liu, Charles Nicholas, and Ting Chen, “Techniques for Gigabyte-Scale N-gram Based Information Retrieval on Personal Computers,” *Proceedings of the 1999 International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA '99)*, Las Vegas, NV, June 1999, pages 1410–1416.
- C25. Michael Shapiro and **Ethan Miller**, “Managing Databases with Binary Large Objects,” *Proceedings of the 16th IEEE Mass Storage System Symposium*, San Diego, CA: IEEE, March 1999, pages 185–193.
- C24. Timothy Gibson, **Ethan L. Miller** and Darrell D. E. Long. “Long-term File Activity and Interference Patterns,” *Proceedings of the 24th Annual International Conference on Computer Measurement and Performance (CMG '98)*, Anaheim, CA: CMG, December 1998, pages 976–987.
- C23. Kennedy Akala, **Ethan Miller**, and Jeff Hollingsworth, “Using Content-Derived Names for Package Management in Tcl,” *Proceedings of the 6th Annual Tcl/Tk Conference*, San Diego, CA: Usenix, September 1998, pages 171–179.
- C22. Theodore Johnson and **Ethan Miller**, “Performance Measurements of Tertiary Storage Devices,” *Proceedings of the 1998 Conference on Very Large Databases (VLDB '98)*, New York, NY: VLDB Foundation, August 1998, pages 50–61.
- C21. Jem Y. Fan, Xiangjun Zhao, J. P. Zhang, Fow-Sen Choa, Yanjie Chai, Jye-Hong Chen, **Ethan Miller**, Howard Motteler, Pao-Lo Liu, Tawee Tanbun-Ek, Patrick Wisk, Won-Tien Tsang, George J. Zyzdik, and Charles A. Burrus, “Wavelength-division-multiplexed (WDM) data block switching for parallel computing and interconnect,” *SPIE International Conference on Applications of Photonic Technology*, Ottawa, Canada: SPIE, July 1998, vol. 3491, pages 634–638.
- C20. Steven Gribble, Gurmeet Singh Manku, Drew Roselli, Eric Brewer, Timothy Gibson, and **Ethan Miller**, “Self-Similarity in File Systems,” *Proceedings of the SIGMETRICS '98 / PERFORMANCE '98 Joint International Conference on Measurement and Modeling of Computer Systems*, Madison, WI: ACM, June 1998, pages 141–150.
- C19. Timothy J. Gibson and **Ethan Miller**, “Long-Term File Activity Patterns in a UNIX Workstation Environment,” *Proceedings of the 6th Goddard Conference on Mass Storage Systems and Technologies / 15th IEEE Symposium on Mass Storage Systems*, College Park, MD: IEEE, March 1998, pages 355–372.
- C18. Theodore Johnson and **Ethan Miller**, “Benchmarking Tape System Performance,” *Proceedings of the 6th Goddard Conference on Mass Storage Systems and Technologies / 15th IEEE Symposium on Mass Storage Systems*, College Park, MD: IEEE, March 1998, pages 95–112.
- C17. R. Scott Cost, Jeegar Lakhani, Ian Soboroff, Tim Finin, **Ethan Miller**, and Charles Nicholas, “TKQML: A Scripting Tool for Building Agents,” *Proceedings of the 1997 Conference on Agent Theories and Agent Languages (ATAL97)*, Newport, RI: AAAI, July 1997, pages 339–343.

- C16. R. Scott Cost, Jeegar Lakhani, Ian Soboroff, Tim Finin, **Ethan Miller**, and Charles Nicholas, “Agent Development Support for Tcl,” *5th Annual Tcl/Tk Workshop '97*, Boston, MA: Usenix, July 1997, pages 177–178.
- C15. Jeff Hollingsworth and **Ethan Miller**, “Using Content-Derived Names for Configuration Management,” *Proceedings of the 1997 Symposium on Software Reusability (SSR '97)*, Boston, MA: IEEE, May 1997, pages 104–109.
- C14. David Ebert, Chris Shaw, Amen Zwa, **Ethan Miller**, and D. A. Roberts, “Interactive Volumetric Information Visualization for Document Corpus Management,” *Proceedings of Graphics Interface*, Kelowna, BC, Canada: Canadian Human-Computer Communications Society, May 1997, pages 121–128.
- C13. Timothy Gibson and **Ethan Miller**, “The Case for Personal Computers as Workstations,” *Proceedings of the 22nd Annual International Conference on Computer Measurement and Performance (CMG '96)*, San Diego, CA: CMG, December 1996, pages 644–652.
- C12. Amen Zwa, David Ebert, and **Ethan Miller**, “Multiresolution Document Analysis with Wavelets,” *Proceedings of the 1996 Conference on Information and Knowledge Management, Workshop on New Paradigms in Information Visualization and Manipulation*, Rockville, MD: ACM, December 1996, pages 50–53.
- C11. David Ebert, Chris Shaw, Amen Zwa, and **Ethan Miller**, “Minimally-immersive Interactive Volumetric Information Visualization,” *Proceedings of IEEE Information Visualization '96*, San Francisco, CA: IEEE, October 1996, pages 66–68.
- C10. **Ethan Miller**, “Towards Scalable Benchmarks for Mass Storage Systems,” *5th NASA Goddard Space Flight Center Conference on Mass Storage Systems and Technologies*, College Park, MD: IEEE & NASA, September 1996, pages 515–528.
- C9. **Ethan Miller** and Randy Katz, “RAMA: Easy Access to a High-Bandwidth Massively Parallel File System,” *Proceedings of the Winter 1995 USENIX Conference*, New Orleans, LA: Usenix, January 1995, pages 59–70.
- C8. Ann Drapeau, Peter Chen, John Hartman, Edward Lee, **Ethan Miller**, Ken Shirriff, Srinu Seshan, Randy Katz, Garth Gibson, and David Patterson, “RAID-II: A Scalable Storage Architecture for High-Bandwidth Network File Service,” *Proceedings of the 21st International Symposium on Computer Architecture*, Chicago, IL: ACM, April 1994, pages 234–244.
- C7. Peter Chen, Edward Lee, Ann Drapeau, Ken Lutz, **Ethan Miller**, Srinu Seshan, Ken Shirriff, David Patterson, and Randy Katz, “Performance and Design Evaluation of the RAID-II Storage Server,” *Proceedings of the International Parallel Processing Symposium Workshop on I/O in Parallel Computer Systems*, Newport Beach, CA: IEEE, April 1993, pages 110–120.
- C6. **Ethan Miller** and Randy Katz, “RAMA: A File System for Massively Parallel Computers,” *Digest of Papers, 12th IEEE Symposium on Mass Storage Systems*, Monterey, CA: IEEE, April 1993, pages 163–168.
- C5. Randy Katz, Peter Chen, Ann Drapeau, Edward Lee, **Ethan Miller**, Srinu Seshan, and David Patterson, “RAID-II: Design and Implementation of a Large Scale Disk Array Controller,” *Proceedings of the VLSI System Design Conference*, Seattle, WA: IEEE, March 1993.

- C4. **Ethan Miller** and Randy Katz, “An Analysis of File Migration in a UNIX Supercomputing Environment,” *Proceedings of the Winter 1993 USENIX Conference*, San Diego, CA: Usenix, January 1993, pages 421–433.
- C3. Randy H. Katz, David A. Patterson, Ann Chervenak-Drapeau, Joel Fine, and **Ethan Miller**, “An Approach to Cost-Effective Terabyte Memory Systems,” *Digest of Papers, Comcon Spring '92, 37th IEEE Computer Society International Conference*, San Francisco, CA: IEEE, February 1992, pages 395–400.
- C2. **Ethan Miller** and Randy Katz, “Input/Output Behavior of Supercomputing Applications,” *Proceedings of Supercomputing '91*, Albuquerque, NM: IEEE, November 1991, pages 567–576.
- C1. **Ethan Miller** and Randy Katz, “Analyzing the I/O Behavior of Supercomputing Applications,” *Digest of Papers, 11th IEEE Symposium on Mass Storage Systems*, Monterey, CA: IEEE, October 1991, pages 51–55.

Chapters in Books

- CH1. Claudia Pearce and **Ethan Miller**, “The TELLTALE Dynamic Hypertext Environment: Approaches to Scalability,” in *Advances in Intelligent Hypertext*, J. Mayfield and C. Nicholas, eds. *Lecture Notes in Computer Science*, Springer-Verlag, October 1997, pages 109–130.

Invited Articles

- I4. Mark W. Storer, Kevin M. Greenan, **Ethan L. Miller**, and Kaladhar Voruganti, “Pergamum: energy-efficient archival storage with disk instead of tape,” *login*: **33**(3), USENIX Association, June 2008, pages 15–21.
- I3. **Ethan L. Miller** and Darrell D. E. Long, “Including Experimental Methods in Operating Systems Courses,” ACM Workshop on Experimental Computer Science (ExpCS 2007), San Diego, CA, June 2007.
- I2. **Ethan L. Miller**, “Dealing with Long-Lived Data in High Performance Storage Systems,” *Storage on the Lunatic Fringe: Beyond Peta-Scale Storage Systems*, workshop at Supercomputing 2003, Phoenix, AZ, November 2003.
- I1. **Ethan Miller**, Darrell Long, William Freeman, and Benjamin Reed, “Strong Security for Distributed File Systems,” *Proceedings of the 20th IEEE International Performance, Computing, and Communications Conference (IPCCC 2001)*, Phoenix, AZ: IEEE, April 2001, pages 34–40.

Work-in-Progress Posters and Presentations

- WP4. Sasha Ames, Carlos Maltzahn, and **Ethan L. Miller**, “A File System Query Language,” poster at the 21st Symposium on Operating Systems Principles (SOSP 2007), Stevenson, WA, October 2007.
- WP3. Mark W. Storer, Kevin Greenan, and **Ethan L. Miller**, “Secure Archival Storage with POTSHARDS,” work-in-progress presentation at the 5th Conference on File and Storage Technologies, San Jose, CA, February 2007.
- WP2. Kevin Greenan and **Ethan L. Miller**, “PRIMS: Making NVRAM Suitable for Extremely Reliable Storage,” work-in-progress presentation at the 5th Conference on File and Storage Technologies, San Jose, CA, February 2007.

- WP1. Andrew Leung and **Ethan L. Miller**, “Scaling Security for Big, Parallel File Systems,” work-in-progress presentation at the 5th Conference on File and Storage Technologies, San Jose, CA, February 2007.

Technical Reports

- T12. Sasha Ames, Carlos Maltzahn, **Ethan L. Miller**, “Quasar: A Scalable Naming Language for Very Large File Collections,” Technical Report UCSC-SSRC-08-04, October 2008.
- T11. Sasha Ames, Carlos Maltzahn, **Ethan L. Miller**, “QUASAR: Interaction with File Systems Using a Query and Naming Language,” Technical Report UCSC-SSRC-08-03, September 2008.
- T10. Andrew W. Leung, Minglong Shao, Timothy Bisson, Shankar Pasupathy, **Ethan L. Miller**, “Spyglass: Fast, Scalable Metadata Search for Large-Scale Storage Systems,” Technical Report UCSC-SSRC-08-01, Storage Systems Research Center, University of California, Santa Cruz, May 2008.
- T9. Kevin M. Greenan, **Ethan L. Miller**, Thomas J. E. Schwarz, S. J., “Analysis and Construction of Galois Fields for Efficient Storage Reliability,” Technical Report UCSC-SSRC-07-09, Storage Systems Research Center, University of California, Santa Cruz, August 2007.
- T8. Mark W. Storer, Kevin Greenan, **Ethan L. Miller**, and Kaladhar Voruganti, “POTSHARDS: Secure Long-Term Archival Storage Without Encryption,” Technical Report UCSC-SSRC-06-03, Storage Systems Research Center, University of California, Santa Cruz, September 2006.
- T7. Sage A. Weil, Scott A. Brandt, **Ethan L. Miller**, and Carlos Maltzahn, “CRUSH: Controlled, Scalable And Decentralized Placement Of Replicated Data,” Technical Report SSRC-06-02, Storage Systems Research Center, University of California, Santa Cruz, March 2006.
- T6. Sage A. Weil, Feng Wang, Qin Xin, Scott A. Brandt, **Ethan L. Miller**, Darrell D. E. Long, and Carlos Maltzahn, “Ceph: a Scalable Object-Based Storage System,” Technical Report SSRC-06-01, Storage Systems Research Center, University of California, Santa Cruz, March 2006.
- T5. Geoff Kuenning and **Ethan L. Miller**, “Anonymization Techniques for URLs and Filenames,” Technical Report UCSC-CRL-03-05, Storage Systems Research Center, University of California, Santa Cruz, September 2003.
- T4. Nathan K. Edel, **Ethan L. Miller**, Karl S. Brandt, and Scott A. Brandt, “Measuring the Compressibility of Metadata and Small Files for Disk/NVRAM Hybrid Storage Systems,” Technical Report UCSC-CRL-03-04, Storage Systems Research Center, University of California, Santa Cruz, July 2003.
- T3. Timothy J. Gibson and **Ethan L. Miller**, “Long-Term File Activity in Diverse UNIX Environments,” Technical Report TR-CS-97-07, University of Maryland Baltimore County, October 1997.
- T2. R. Scott Cost, Ian Soboroff, Jeegar Lakhani, Tim Finin, **Ethan Miller**, and Charles Nicholas, “TKQML: A KQML Extension to Tcl,” Technical Report TR-CS-97-04, University of Maryland Baltimore County, July 1996.
- T1. **Ethan Miller** and Jeffrey Hollingsworth, “Using Content-Derived Names for Caching and Software Distribution,” Technical Report TR-CS-96-08, University of Maryland Baltimore County, July 1996. Also available as UMIACS Technical Report TR-96-55.

Miscellanea

- M1. Randy Katz, John Ousterhout, David Patterson, Peter Chen, Ann Chervenak, Rich Drewes, Garth Gibson, Edward Lee, Ken Lutz, Ethan Miller, and Mendel Rosenblum, “A Project on High-Performance I/O Subsystems,” *Computer Architecture News* **17**(5):24–31, September 1989.

SOFTWARE

- SW1. GaloisField: A high-performance Galois field library written in C/C++ with an efficient Python interface. This software was co-authored with Kevin Greenan, one of my Ph. D. students.
- SW2. DLX0S: A CPU simulator and operating system used in undergraduate operating systems classes at the University of California, Santa Cruz, the University of Arizona, Purdue, the University of Maryland Baltimore County, and elsewhere.
- SW3. esim: A simple digital logic design language and simulator. Used for teaching computer architecture at the University of Maryland Baltimore County and elsewhere.

PROFESSIONAL ACTIVITIES

Service to Professional Societies

- 2001– Chair, *IEEE Technical Committee on Operating Systems and Applications Environments*.
- 1996– Member, Executive Committee, *IEEE Technical Committee on Mass Storage Systems*.
- 1995–2000 University Liaison for the Usenix Association.

Conference Organization

- 2009 Program Committee: 7th File and Storage Technologies Conference (FAST 2009); 17th Conference on Measurement and Simulation of Computer and Telecommunication Systems (MASCOTS 2006); The Israeli Experimental Systems Conference (SYSTOR 2009).
- 2008 Program co-Chair, 16th Conference on Measurement and Simulation of Computer and Telecommunication Systems (MASCOTS 2008).
Program Committee: 28th International Conference on Distributed Computing Systems (ICDCS 2008); 4th ACM Workshop on Storage Survivability and Security (StorageSS 2008); 5th IEEE International Workshop on Storage Network Architecture and Parallel I/O (SNAPI); 1st International Workshop on Storage and I/O Virtualization, Performance, Energy, Evaluation and Dependability (SPEED 2008); 3rd International Workshop on Software Support for Portable Storage (IWSSPS 2008).
- 2007 Program Chair, 24th IEEE Conference on Mass Storage Systems and Technologies (MSST 2007).
Area Vice-Chair (Data-Intensive Computing), 36th International Conference on Parallel Processing (ICPP 2007).
Program Committee: 3rd ACM Workshop on Storage Survivability and Security (StorageSS 2007); 4th International IEEE Security in Storage Workshop (SISW 2007).
- 2006 Program co-Chair, 2nd ACM Workshop on Storage Survivability and Security (StorageSS 2006).
Program Committee: 23rd IEEE / 14th NASA Goddard Conference on Mass Storage Systems and Technologies (MSST 2006); 14th Conference on Measurement and Simulation of Computer and Telecommunication Systems (MASCOTS 2006); 7th Workshop on Distributed Data and Structures (WDAS 2006); 2nd International Workshop on Software Support for Portable Storage (IWSSPS 2006).

- 2005 Program Committee: 13th Conference on Measurement and Simulation of Computer and Telecommunication Systems (MASCOTS 2005); 22nd IEEE / 13th NASA Goddard Conference on Mass Storage Systems and Technologies (MSST 2005); International Symposium on Emergence of Globally Distributed Data (publications chair); International Workshop on Software Support for Portable Storage (IWSSPS 2005); ACM Workshop on Storage Survivability and Security (StorageSS); 3rd International IEEE Security in Storage Workshop; Invited Talks co-chair, 2005 USENIX Technical Conference.
- 2004 Program Committee: 21st IEEE / 12th NASA Goddard Conference on Mass Storage Systems and Technologies (MSST 2004); 2nd Intelligent Storage Workshop; 6th Workshop on Distributed Data and Systems (WDAS 2004); 12th Conference on Measurement and Simulation of Computer and Telecommunication Systems (MASCOTS 2004), publications chair; Supercomputing (SC) 2004.
- 2003 Program Committee: 20th IEEE / 11th NASA Goddard Conference on Mass Storage Systems and Technologies (MSST 2003), publications chair; 2nd International IEEE Security in Storage Workshop.
Steering Committee: HotOS-IX: Ninth Workshop on Hot Topics in Operating Systems.
- 2002 Program Committee: 1st File and Storage Technologies Conference (FAST 2002), publications chair; 19th IEEE Mass Storage System Symposium / 10th NASA Goddard Mass Storage and Technologies Conference; 1st International IEEE Security in Storage Workshop; IEEE International Performance, Computing, and Communications Conference.
- 2001 Program Committee: 18th IEEE Mass Storage System Symposium / 9th NASA Goddard Mass Storage and Technologies Conference.
- 2000 Program Committee: 17th IEEE Mass Storage System Symposium / 8th NASA Goddard Mass Storage and Technologies Conference.
- 1999 Program Committee: 16th IEEE Mass Storage System Symposium / 7th NASA Goddard Mass Storage and Technologies Conference (publications chair).
- 1998 Program Committee: 15th IEEE Mass Storage System Symposium / 6th NASA Goddard Mass Storage and Technologies Conference.
- 1996 Program Committee: Workshop on I/O in Parallel and Distributed Systems (IOPADS).

Reviewer of Technical Papers and Proposals

- 2008 *IEEE Transactions on Parallel and Distributed Systems*; *ETRI (Electronics and Telecommunications Research Institute) Journal*; 6th Conference on File Systems and Technologies (FAST); National Science Foundation (CISE panel).
- 2007 *IEEE Transactions on Parallel and Distributed Systems*; 5th Conference on File Systems and Technologies (FAST); WCNC; National Science Foundation (CISE panel); National Science Foundation (CAREER panel); UC Discovery Grant.
- 2006 Conference on Dependable Systems and Networks (DSN); 7th Conference on Operating Systems Design and Implementation (OSDI); National Science Foundation (CAREER panel); National Science Foundation (SBIR panels).
- 2005 *IEEE Transactions on Computers*; 4th Conference on File and Storage Technologies (FAST); 19th International Symposium on Distributed Computing (DiSC 2005); National Science Foundation (CISE panel).
- 2004 3rd Conference on File and Storage Technologies (FAST); Conference on Dependable Systems and Networks (DSN); 6th Symposium on Operating Systems Design and Implementa-

- tion (OSDI); *Software—Practice and Experience*; *ACM Transactions on Computer Systems*; National Science Foundation (CISE panel).
- 2003 *IEEE Transactions on Parallel and Distributed Systems*; Usenix Technical Conference; 2nd Conference on File and Storage Technologies; SC2003; National Science Foundation (SBIR panel).
- 2002 *IEEE Transactions on Computers*, *IEEE Transactions on Parallel and Distributed Systems*, Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), 5th Symposium on Operating Systems Design and Implementation (OSDI), Usenix Annual Technical Conference, National Science Foundation (SBIR panel).
- 2001 Usenix Annual Technical Conference; *IEEE Transactions on Computers*; National Science Foundation (ITR panel, SBIR panel).
- 1990–2000 National Science Foundation (panel reviewer); *ACM Transactions on Computer Systems*; *IEEE Computer*; International Symposium on Computer Architecture; *Journal of Parallel and Distributed Computing*; Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS); Symposium on Operating Systems Design and Implementation (OSDI); USENIX Annual Technical Conference.

Reviewer for Publishers

- 1997 Oxford University Press

Membership in Professional Associations

- Senior Member, IEEE Computer Society.
 Member, Association for Computing Machinery.
 Member, Usenix Association.
 Member, Sigma Xi.

Invited Talks

- 2008 “Building Reliable, Efficient Metadata Storage on NVRAM,” Hanyang University, Seoul, Korea.
- 2008 “Building Reliable, Efficient Metadata Storage on NVRAM,” Samsung, Seoul, Korea.
- 2008 “Reliable and Efficient Metadata Storage and Indexing Using NVRAM,” invited talk at the International Workshop on Large-Scale NVRAM Technology (NVRAMOS08) Workshop, Jeju, Korea.
- 2008 “Pergamum: Building Evolvable, Reliable, Searchable, Energy-Efficient Petascale Archival Storage from Disks,” Yahoo!, Santa Clara, CA.
- 2008 “Scaling Metadata and Indexing for Petabyte-Scale Storage,” invited talk at the 5th IEEE International Workshop on Storage Network Architecture and Parallel I/O, Baltimore, MD.
- 2008 Panel member, “SSD in the Enterprise,” Flash Memory Summit, Santa Clara, CA.
- 2008 “Highly Scalable Metadata Search and Indexing,” Metadata Research Panel, HEC FSIO 2008 Conference, Arlington, VA.
- 2008 “Search and indexing for petabyte-scale storage and beyond,” *Overcoming I/O Bottlenecks in Full Data Path Processing: Intelligent, Scalable Data Management from Data Ingest to Computation Enabling Access and Discovery*, National Science Foundation workshop, Arlington, VA.

- 2008 “Pergamum: Replacing Tape with Energy Efficient, Reliable, Disk-Based Archival Storage,” Association of Moving Image Archivists, Los Angeles, CA.
- 2008 “Pergamum: Evolvable Energy Efficient, Reliable, Disk-Based Archival Storage,” IBM off-site research review, San Jose, CA.
- 2008 “Pergamum: Replacing Tape with Energy Efficient, Reliable, Disk-Based Archival Storage,” LSI Logic, Boulder, CO.
- 2008 “Pergamum: Replacing Tape with Energy Efficient, Reliable, Disk-Based Archival Storage,” 8th Intelligent Storage Workshop, University of Minnesota, Minneapolis, MN.
- 2008 “Pergamum: Replacing Tape with Energy Efficient, Reliable, Disk-Based Archival Storage,” Symantec, Mountain View, CA.
- 2008 “Ceph: A Scalable, High-Performance Distributed File System,” University of Paderborn, Paderborn, Germany.
- 2008 “Pergamum: Replacing Tape with Energy Efficient, Reliable, Disk-Based Archival Storage,” University of Paderborn, Paderborn, Germany.
- 2008 “Scalable Security in Ceph,” Winter 2008 SNIA Security Summit, San Jose, CA.
- 2007 “Recent Research at the Storage Systems Research Center,” Data Domain, Sunnyvale, CA.
- 2007 “POTSHARDS: Secure Long-Term Archival Storage Without Encryption,” University of California, Santa Barbara.
- 2007 “POTSHARDS: Secure Long-Term Archival Storage Without Encryption,” TRUST seminar, University of California, Berkeley.
- 2007 “Disaster Tolerance Codes,” Agami, Sunnyvale, CA.
- 2007 “POTSHARDS: Secure Long-Term Archival Storage Without Encryption,” VMware, Palo Alto, CA.
- 2007 “Reliability Mechanisms for File Systems Using Non-Volatile Memory as a Metadata Store,” Network Appliance, Sunnyvale, CA.
- 2007 “Ceph: A Scalable, High-Performance Distributed File System,” National Technology Alliance Data Retention and Processing Workshop, Herndon, VA.
- 2007 Visited the Data Storage Institute, affiliated with the National University of Singapore. Gave five talks at DSI and discussed DSI-UCSC collaborations with DSI researchers:
- “Ceph: A Scalable, High-Performance Distributed File System”
 - “Maat: Scalable Security for High Performance, Petascale Storage”
 - “POTSHARDS: Secure Long-Term Archival Storage Without Encryption”
 - “Store, Forget, and Check: Using Algebraic Signatures to Check Remotely Administered Storage”
 - “Improving File System Performance and Reliability with Non-Volatile Memory”
- 2007 Panel member, Education Roundtable, Workshop on Experimental Computer Science (EX-PCS 2007).
- 2007 Panel member, “High Performance I/O: The Road Less Traveled?” 8th LCI International Conference on High-Performance Clustered Computing.
- 2007 “Using Massive Arrays of Idle Disks (MAIDs) for Long-Term Storage in Digital Libraries,” Digital Library Forum, Pasadena, CA.
- 2007 “POTSHARDS: Secure Long-Term Archival Storage Without Encryption,” University of New Mexico, Albuquerque, NM.

- 2007 “Ceph: Scalable, Reliable, Secure, High-Performance Storage,” Aster Data Systems, Redwood Shores, CA.
- 2007 “Improving File System Performance and Reliability with Non-Volatile Memory,” eBay, San Jose, CA.
- 2007 “POTSHARDS: Secure Long-Term Archival Storage Without Encryption,” Winter 2007 SNIA Storage Security Workshop, San Diego, CA.
- 2007 “Ceph: Scalable, Reliable, Secure, High-Performance Storage,” University of New Mexico, Albuquerque, NM.
- 2007 “POTSHARDS: Secure Long-Term Archival Storage Without Encryption,” IBM Almaden Research Lab, San Jose, CA.
- 2006 “POTSHARDS: Secure Long-Term Archival Storage Without Encryption,” Data Domain, Santa Clara, CA.
- 2006 “Ceph: A Scalable, High-Performance Distributed File System,” University of Rhode Island, Kingston, RI.
- 2006 “POTSHARDS: Secure Long-Term Archival Storage Without Encryption,” Brown University, Providence, RI.
- 2006 “Store, Forget, and Check: Using Algebraic Signatures to Check Remotely Administered Storage,” Symantec, Mountain View, CA.
- 2006 “Improving File System Usability, Performance and Reliability with Magnetic RAM,” Los Alamos National Laboratory, Los Alamos, NM.
- 2006 “Security in Large Scale File Systems,” Sandia National Laboratory, Albuquerque, NM.
- 2006 “Security in Large Scale File Systems,” National Security Agency, Fort Meade, MD.
- 2006 “POTSHARDS: Secure Long-Term Archival Storage Without Encryption,” IBM Storage Strategy Review (poster).
- 2006 “OSD Security,” Winter 2006 SNIA Storage Security Workshop, San Diego, CA.
- 2005 “Using Object Storage Devices to Provide Highly Reliable, Highly Scalable Storage,” keynote address at IBM Haifa Workshop on Systems and Storage.
- 2005 “Peta-scale Data Storage with Commodity Components,” Stony Brook University, Stony Brook, NY.
- 2005 “Peta-scale Data Storage with Commodity Components,” University of Maryland Baltimore County.
- 2005 “Object-Based Storage Systems,” Yahoo.
- 2005 “Ensuring Reliability for Petabyte-Scale Storage Systems,” Hewlett Packard Laboratories, Palo Alto, CA.
- 2004 “Is Massive Election Computer Fraud Possible?” panel discussion on *Voices*, Santa Cruz Community Television.
- 2004 “Dynamic Metadata Management for Large Distributed Storage Clusters,” Sandia National Laboratory, Albuquerque, NM.
- 2003 Panel member, “Roadmaps, Research Thrusts, and More Visibility For Storage Issues,” *Storage on the Lunatic Fringe: Beyond Peta-Scale Storage Systems*, Supercomputing 2003 workshop.
- 2003 “Storage Research in the UCSC Storage Systems Research Center,” Veritas Corporation, Mountain View, CA.

- 2003 “RUSH: A Family of Algorithms for the Placement of Objects in Distributed Object Storage Devices,” Hewlett Packard Laboratories, Palo Alto, CA.
- 2003 Panel member, “Emerging Object/Active-Storage Technologies,” 20th IEEE / 11th NASA Goddard Conference on Mass Storage Systems and Technologies.
- 2001 “Improving File System Performance Using Magnetic RAM (MRAM),” Hewlett Packard Laboratories.
- 2001 “Strong Security for Network-Attached Storage,” IBM Research.
- 2000 “RAMA: A Reliable High-Performance Scalable File System,” Army Research Lab.
- 2000 “RAMA: A Reliable High-Performance Scalable File System,” University of California, Santa Cruz.
- 1999 “Building Large Scalable File Systems,” Lawrence Livermore National Laboratory.
- 1999 “Scaling TELLTALE from Megabytes to Gigabytes,” Department of Defense.
- 1998 “Content-Derived Names: Security, Distribution, and Version Management,” University of Wisconsin, Milwaukee.
- 1997 “Performance Evaluation and Estimation,” RWD, Inc. (Columbia, MD).
- 1997 “Challenges in Mass Storage,” Silicon Graphics.
- 1996 “RAMA: An Easy-to-Use, Scalable High Performance Parallel File System,” Sandia National Laboratory.

UNIVERSITY SERVICE

Departmental and School Service

- 2008–09 Chair, School of Engineering Computing Infrastructure Committee; Computer Science Dept. Personnel Committee; Computer Science Dept. Computing Committee; Computer Science Dept. Seminar Committee.
- 2007–08 Chair, School of Engineering Computing Infrastructure Committee; Computer Science Dept. Personnel Committee; Computer Science Dept. Computing Committee.
- 2006–07 School of Engineering Computing Infrastructure Committee (member, Fall; chair, Spring); Computer Science Dept. Personnel Committee; Computer Science Dept. Computing Committee.
- 2005–06 School of Engineering Computing Infrastructure Committee (chair, Fall; member Winter and Spring); Computer Science Dept. Personnel Committee; Computer Science Dept. Computing Committee.
- 2004–05 Chair, School of Engineering Computing Infrastructure Committee; Computer Science Dept. Faculty Recruiting Committee; Computer Science Dept. Personnel Committee (chaired 2 reviews); Computer Science Dept. Computing Committee.
- 2003–04 Chair, School of Engineering Computing Infrastructure Committee; Chair, Computer Science Dept. Computing Committee; Computer Science Dept. Faculty Recruiting Committee.
- 2002–03 Chair, School of Engineering Computing Infrastructure Committee.
- 2001–02 School of Engineering Computing Infrastructure Committee; Computer Science Dept. Faculty Recruiting Committee.
- 2000–01 Computer Science Dept. Faculty Recruiting Committee.

University Committees

- 2007–08 Committee on Faculty Research Lecture
 2006–07 Committee on Computing and Telecommunications (Fall & Spring)
 2005–06 Chair, Committee on Computing and Telecommunications.
 2004–05 Chair, Committee on Computing and Telecommunications.
 2003–04 Committee on Computing and Telecommunications (Winter & Spring).
 2003 Information Technology Vision Committee.

System-Wide Committees

- 2005–06 Committee on Information Technology and Telecommunications Policy.
 2004–05 Committee on Information Technology and Telecommunications Policy.

COMMUNITY SERVICE

- 2003–05 Member, Santa Cruz Hillel Board of Directors
 2002– Alumni interviewer, Brown University

ADVISING**Advisor to Continuing Graduate Students**

- | | |
|---------------------|------------------------|
| Ian Adams | M. S. |
| Kevin Greenan | Ph. D. (expected 2009) |
| Keren Jin | Ph. D. |
| Andrew Leung | Ph. D. (expected 2009) |
| Aleatha Parker-Wood | Ph. D. |
| Mark Storer | Ph. D. (expected 2008) |
| Avani Wildani | Ph. D. |

Masters of Science Advisor

- | | | |
|------|-------------------------|--|
| 2008 | Casey Marshall | non-thesis |
| 2008 | Danni Fu | non-thesis |
| 2007 | Jeff Hagen | <i>Ladon: A Framework for Peer-to-Peer Backup</i> |
| 2007 | Max Mehech | non-thesis |
| 2007 | Andrew Leung | <i>Security in Scalable Storage Systems</i> |
| 2006 | Kevin Greenan | non-thesis |
| 2006 | Mark Storer | non-thesis |
| 2005 | Christopher Olson | non-thesis |
| 2005 | Chengyu Sung | <i>Integrating Pictorial Identity into Secure Email</i> |
| 2005 | Sasha Ames | non-thesis |
| 2004 | R. J. Honicky | <i>Object Placement Algorithms for OBSD Systems</i> |
| 2001 | Kennedy Akala | non-thesis |
| 2000 | Vivekand Krishnamoorthi | <i>A Comparison of Long Term File Migration Algorithms</i> |
| 1999 | Ting Chen | non-thesis |
| 1998 | Junli Liu | non-thesis |
| 1998 | Michael Shapiro | non-thesis |
| 1998 | Dan Shen | <i>Experiments with Large-Scale N-gram Based Information Retrieval</i> |
| 1998 | Changgong Zhang | non-thesis |
| 1997 | Eric Robertson | non-thesis |
| 1996 | Arun C. Mahendran | non-thesis |
| 1996 | Amen Zwa | non-thesis |

Masters of Science Reading Committee Member

2007	Corrie Scalisi	<i>Prof. Manfred Warmuth</i>
2006	Suma Potluri	<i>Prof. E. James Whitehead</i>
2006	Nikhil Bobb	<i>Prof. Scott Brandt</i>
2005	Travis Odegaard	<i>Prof. Scott Brandt</i>
2004	Svetlana Kagan	<i>Prof. Darrell Long</i>
2004	Deepa Tuteja (thesis)	<i>Prof. Scott Brandt</i>
2003	Suruchi Malapture (thesis)	<i>Prof. Scott Brandt</i>
2003	Karen Glocer	<i>Prof. Darrell Long</i>
2003	Ravindra Vaishampayan	<i>Prof. J. J. Garcia-Luna</i>
2003	Caixue Lin	<i>Prof. Scott Brandt</i>
1999	Amy Germida	<i>Prof. James Plusquellic (UMBC)</i>
1996	Greg Sylvain	<i>Prof. Tim Finin (UMBC)</i>
1996	Chetan Shah	<i>Prof. Charles Nicholas (UMBC)</i>
1996	Scott Stewart	<i>Prof. Tim Finin (UMBC)</i>

Doctoral Advisor

2005	Qin Xin	<i>Understanding and Coping with Failures in Large-Scale Storage Systems</i>
2004	Ismail Ari	<i>Design and Management of Globally-Distributed Network Caches</i>
2002	Naomi Avigdor	<i>Building a Scalable and Reliable Parallel File System Using Commodity Computers</i>
2000	William Freeman	<i>Decentralized Security for Network-Attached Storage</i>
1998	Timothy Gibson	<i>Long-term UNIX File System Activity and the Efficacy of Automatic File Migration</i>

Doctoral Dissertation Reading Committee Member

2008	Vinay Pai	<i>Prof. Erez Zadok (Stony Brook University)</i>
2008	Guozheng Ge	<i>Prof. E. James Whitehead</i>
2007	Sage Weil	<i>Prof. Scott Brandt</i>
2006	Nikolai Joukov	<i>Prof. Erez Zadok (Stony Brook University)</i>
2006	Lawrence You	<i>Prof. Darrell Long</i>
2006	Feng Wang	<i>Prof. Scott Brandt</i>
2005	Bo Hong	<i>Profs. Darrell Long and Scott Brandt</i>
2005	Scott Banachowski	<i>Prof. Scott Brandt</i>
2002	Ahmed Amer	<i>Prof. Darrell Long</i>
2002	Tsozen (Frank) Yeh	<i>Prof. Darrell Long</i>
2000	Ian Soboroff	<i>Prof. Charles Nicholas (UMBC)</i>

Doctoral Qualifying Exam Committee Member

2008	Dhananjay Sampath	<i>Prof. J. J. Garcia-Luna</i>
2008	Kevin Greenan	<i>Prof. Ethan Miller</i>
2008	Neerja Bhatnagar	<i>Prof. Ethan Miller</i>
2007	Mark W. Storer	<i>Prof. Ethan Miller</i>
2007	Deepavali Bhagwat*	<i>Prof. Darrell Long</i>
2006	Sage Weil*	<i>Prof. Scott Brandt</i>
2006	Elias Sinderson	<i>Prof. E. James Whitehead</i>
2006	Nikolai Joukov	<i>Prof. Erez Zadok (Stony Brook University)</i>
2005	David Pease	<i>Prof. Darrell Long</i>
2004	Guozheng Ge*	<i>Prof. E. James Whitehead</i>
2004	Damian Cieslicki	<i>Prof. Thomas Schwarz (Santa Clara University)</i>
2003	Scott Banachowski	<i>Prof. Scott Brandt</i>
2003	Qin Xin	<i>Prof. Ethan Miller</i>
2003	Bo Hong	<i>Prof. Darrell Long</i>
2003	Feng Wang	<i>Prof. Scott Brandt</i>
2002	Ismail Ari	<i>Prof. Ethan Miller</i>
2001	Tsozen Yeh	<i>Prof. Darrell Long</i>
1999	Naomi Avigdor	<i>Prof. Ethan Miller (UMBC)</i>
1999	William Freeman	<i>Prof. Ethan Miller (UMBC)</i>
1998	Ian Soboroff	<i>Charles Nicholas (UMBC)</i>
1997	Timothy Gibson	<i>Prof. Ethan Miller (UMBC)</i>
1995	Vincent Marier	<i>Prof. Deepinder Sidhu (UMBC)</i>

Note: * indicates that I served as qualifying exam committee chair.

COURSES TAUGHT

Note that classes from Fall 1994 through Spring 2000 are semester classes, and classes from Fall 2000 onward are quarter classes. Courses in italics are those for which I either developed or significantly revised the curriculum.

Undergraduate

Fall 2008 *Introduction to Programming in Python*
Spring 2008 Distributed Systems
Fall 2007 Operating Systems
Spring 2007 Computer Security
Spring 2006 Operating Systems
Winter 2006 *Distributed Systems*
Winter 2005 Computer Security
Fall 2004 Operating Systems
Spring 2004 Computer Security
Fall 2003 Operating Systems
Spring 2003 *Computer Security*
Winter 2003 Operating Systems
Winter 2002 Introduction to Data Structures
Fall 2001 Operating Systems
Fall 2000 Operating Systems
Fall 1999 Operating Systems
Spring 1999 Operating Systems
Fall 1997 Computer Architecture
Spring 1997 Computer Architecture
Spring 1996 Computer Architecture
Spring 1995 *Computer Architecture*
Fall 1994 Computer Architecture

Graduate

Winter 2008 Storage Systems
Spring 2007 Distributed Systems
Fall 2005 Operating Systems
Spring 2005 Distributed Systems
Winter 2004 Storage Systems
Fall 2002 Operating Systems
Spring 2002 *Distributed Systems*
Spring 2001 Computer Security
Spring 2000 Computer Architecture
Spring 2000 Storage Systems
Fall 1998 Computer Architecture
Spring 1998 Storage Systems

Fall 1997 Computer Architecture
Spring 1997 Operating Systems
Fall 1996 Computer Architecture
Fall 1995 Computer Architecture